

**UNITED STATES DISTRICT COURT
WESTERN DISTRICT OF TEXAS
AUSTIN DIVISION**

Intellectual Ventures I LLC and
Intellectual Ventures II LLC,

Plaintiffs

v.

VMware, Inc.,

Defendant.

Civil Action No. 1:19-cv-01075-ADA

PLAINTIFFS' CLAIM CONSTRUCTION BRIEF

TABLE OF CONTENTS

I.	Introduction	1
II.	Legal Standards	1
	A. Claim Construction Generally	1
	B. Means-Plus-Function Claim Construction Analysis	2
	C. Indefiniteness	2
III.	The Disputed Claim Terms	3
	A. The '752 Patent	3
	B. The '686 and '726 Patents	8
	1. “modify a resource allocation” / “modified resource allocation” / “modify[ing] [the] computer resources allocated to a virtual server” ('686 patent Claims 5, 6, 7)	9
	2. “modify[ing] a resource allocation for the virtual server” / “modifying [the] computer resources allocated to a virtual server” ('726 patent Claims 1, 4, 5 and 8)	9
	3. “resource denials” ('726 patent Claims 1, 4, 5 and 8)	10
	4. “resource unavailable messages” / “denied requests to modify a resource allocation” / “resource unavailable messages resulting from denied requests to modify a resource allocation” ('686 patent Claims 5, 6, 7)	11
	5. “virtual server” ('686 patent Claims 5, 6, 7); ('726 patent Claims 1-11)	12
	6. “quality of service guarantee” ('726 patent Claims 1 and 4)	15
	7. “determining that a second physical host can accommodate the requested modified resource allocation” ('686 patent Claims 5, 6 and 7)	16
	8. “determination that a virtual server is overloaded” ('686 patent Claims 5, 6 and 7) .	18
	9. “virtual server overload signal” ('726 patent Claims 1, 4, 5 and 8)	20
	10. Alleged Means Plus Function Elements ('686 patent Claim 7) ('726 patent Claims 1, 3, 5 and 7)	21
	C. The '818 Patent	29
	1. “virtual [network/storage network] interface layer of an application server” (Claims 1, 17, 32, 42)	29

2. “hierarchical token bucket resource allocation” / “token(s)” (Claims; 1, 17, 30, 32, 33, 37-42).....	31
3. “enforcing . . .” / “receiv[e/ing] . . .” / “classify[ing] . . .” / “compar[e/ing] . . .” / “forward[ing] . . .” / “buffer[ing] . . .” (Claims; 1, 17, 30, 32, 37-39, 42).....	33
4. “maintain[ing] a connection over a network fabric” (Claims 1, 17, 30, 32, 42).....	34
5. Alleged Means Plus Function Elements (Claim 17) (individually set forth in Ex. C)	36
D. The ‘051 Patent	38
1. “customer forwarding [table(s)/information]” (claims: 1, 3).....	38
2. “virtual server(s)” (claims; 1, 3, 6)	41
3. “physical interface(s)” (claims; 1, 3)	42
4. “storing . . .” / “determining . . .” / “using . . .” (claims; 1, 3)	43
IV. Conclusion	47

TABLE OF AUTHORITIES

Cases

<i>Acumed LLC v. Stryker Corp.</i> , 483 F.3d 800 (Fed. Cir. 2007).....	45
<i>Advanced Mktg. Sys., LLC v. CVS Pharmacy, Inc.</i> , 2016 WL 1741396 (E.D. Tex. May 3, 2016)	22
<i>Aloft Media, LLC v. Adobe Sys., Inc.</i> , 570 F. Supp. 2d 887 (E.D. Tex. 2008)	24, 26, 27
<i>Apple Inc. v. Motorola, Inc.</i> , 757 F.3d 1286 (Fed. Cir. 2014).....	37
<i>Bell Commc'ns. Research, Inc. v. Vitalink Commc'ns. Corp.</i> , 55 F.3d 615 (Fed. Cir. 1995).....	40
<i>Catalina Mktg. Int'l, Inc. v. Coolsavings.com, Inc.</i> , 289 F.3d 801 (Fed. Cir. 2002).....	33
<i>Collaborative Agreements, LLC v. Adobe Sys. Inc.</i> , 2015 WL 2250391 (W.D. Tex. May 12, 2015).....	23
<i>Comark Commc'ns, Inc. v. Harris Corp.</i> , 156 F.3d 1182 (Fed. Cir. 1998).....	9, 10, 14
<i>Cont'l Circuits LLC v. Intel Corp.</i> , 915 F.3d 788 (Fed. Cir. 2019).....	31
<i>CSB-Sys. Int'l, Inc. v. SAP America, Inc.</i> , 2011 WL 3240838 (E.D. Pa. July 28, 2011).....	23
<i>Dealertrack, Inc. v. Huber</i> , 674 F.3d 1315 (Fed. Cir. 2012).....	10
<i>Eolas Techs., Inc. v. Adobe Sys., Inc.</i> , 810 F. Supp. 2d 795 (E.D. Tex. 2011)	24, 26, 27
<i>Epistar Corp. v. Int'l Trade Comm'n</i> , 566 F.3d 1321 (Fed. Cir. 2009).....	35, 46
<i>Fisher-Rosemount Sys. v. ABB Ltd.</i> , 2019 WL 6830806 (S.D. Tex. Dec. 12, 2019)	21, 22
<i>Intellectual Ventures II LLC v. Bitco Gen. Ins. Corp.</i> , 2016 WL 125594 (E.D. Tex. Jan. 11, 2016)	24
<i>Inventio AG v. ThyssenKrupp Elevator Am. Corp.</i> , 649 F.3d 1350 (Fed. Cir. 2011).....	36
<i>Liebel-Flarsheim Co. v. Medrad, Inc.</i> , 358 F.3d 898 (Fed. Cir. 2004).....	19
<i>Markman v. Westview Instruments, Inc.</i> , 52 F.3d 967 (Fed. Cir. 1995).....	7
<i>Masco Corp. v. United States</i> , 303 F.3d 1316 (Fed. Cir. 2002).....	2
<i>Maurice Mitchell Innovations, L.P. v. Intel Corp.</i> , 2006 WL 1751779 (E.D. Tex. June 21, 2006).....	3, 4
<i>Micro Chem., Inc. v. Great Plains Chem. Co.</i> , 194 F.3d 1250 (Fed. Cir. 1999).....	7

<i>Nautilus, Inc. v. Biosig Instruments, Inc.</i> , 572 U.S. 898 (2014).....	3
<i>Netfuel, Inc. v. F5 Networks, Inc.</i> , 2017 WL 2834538 (N.D. Ill. June 29, 2017)	22
<i>O2 Micro Int’l Ltd. v. Beyond Innovation Tech. Co.</i> , 521 F.3d 1351 (Fed. Cir. 2008).....	45
<i>Phillips v. AWH Corp.</i> , 415 F.3d 1303 (Fed. Cir. 2005).....	1, 2, 7
<i>Pitney Bowes, Inc. v. Hewlett-Packard Co.</i> , 182 F.3d 1298 (Fed. Cir. 1999).....	10
<i>Retractable Techs., Inc. v. Becton, Dickinson & Co.</i> , 653 F.3d 1296 (Fed. Cir. 2011).....	35, 46
<i>RLIS, Inc. v. Allscripts Healthcare Solutions, Inc.</i> , 2013 WL 3772472 (S.D. Tex. July 16, 2013).....	24, 26, 27
<i>Sonix Tech. Co. v. Publ’ns Int’l, Ltd.</i> , 844 F.3d 1370 (Fed. Cir. 2017).....	3
<i>Sound View Innovations, LLC v. Facebook, Inc.</i> , 2017 WL 2221177 (D. Del. May 19, 2017).....	22
<i>Thorner v. Sony Computer Entm’t Am. LLC</i> , 669 F.3d 1362 (Fed. Cir. 2012).....	2
<i>Vitronics Corp. v. Conceptoronic, Inc.</i> , 90 F.3d 1576 (Fed. Cir. 1996).....	1
<i>Watts v. XL Sys., Inc.</i> , 232 F.3d 877 (Fed. Cir. 2000).....	2, 36
<i>Williamson v. Citrix Online, LLC</i> , 792 F.3d 1339 (Fed. Cir. 2015).....	2, 21, 38
<i>Zeroclick, LLC v. Apple Inc.</i> , 891 F.3d 1003 (Fed. Cir. 2018).....	38
Other Authorities	
IEEE Standard Dictionary of Electrical and Electronics Terms, p. 1182 (6th ed. 1996)	14

Per the Court’s Order Governing Proceedings – Patent Case, plaintiffs, Intellectual Ventures I LLC and Intellectual Ventures II LLC (together “IV”), respectfully submit this claim construction brief in support of their proposed constructions for the disputed terms.

I. Introduction

IV alleges that VMware infringes five patents owned by IV: United States Patent Nos. 7,949,752 (“the ’752 patent”); RE 44,686 (“the ’686 patent”); RE 42,726 (“the ’726 patent”); RE 43,051 (“the ’051 patent”) and RE 44,818 (“the ’818 patent”).¹ IV’s patents teach systems and methods to virtualize data centers and enable the use of cloud computing to deliver computer services. VMware’s products and services deliver virtualization and cloud computing services to its customers.

II. Legal Standards

A. Claim Construction Generally

In resolving claim construction disputes, the Court considers three “intrinsic” sources: (1) the claims; (2) the patent specifications; and (3) the prosecution histories. *Vitronics Corp. v. Conceptronic, Inc.*, 90 F.3d 1576, 1582 (Fed. Cir. 1996). Courts may also rely on extrinsic evidence, such as dictionaries or treatises. *Phillips v. AWH Corp.*, 415 F.3d 1303, 1317 (Fed. Cir. 2005). A trial court is not required to follow any set formula to construe disputed claim terms “[n]or is the court barred from considering any particular sources or required to analyze sources in any specific sequence . . .” *Id.* at 1324.

“The claims of a patent define the invention to which the patentee is entitled the right to exclude.” *Id.* at 1312. “The words of a claim are generally given their ordinary and customary meaning as understood by a person of ordinary skill in the art when read in the context of the

¹ References in this brief to “Ex. __” are to exhibits attached to the Declaration of Jonathan R. DeBlois (“DeBlois Decl.”).

specification and prosecution history.” *Thorner v. Sony Computer Entm’t Am. LLC*, 669 F.3d 1362, 1365 (Fed. Cir. 2012). *See also Phillips*, 415 F.3d at 1315 (“In some cases, the ordinary meaning of the claim language... may be readily apparent even to lay judges, and claim construction in such cases involves little more than the application of the widely accepted meaning of commonly understood words.”). There are only two exceptions to the general rule that claims are given their ordinary and customary meaning: “1) when a patentee sets out a definition and acts as his own lexicographer, or 2) when the patentee disavows the full scope of a claim term either in the specification or during prosecution.” *Thorner*, 669 F.3d at 1365. Extrinsic evidence may be consulted and used to construe claim terms only if such evidence does not contradict the intrinsic evidence. *Phillips*, 415 F.3d at 1322-23.

B. Means-Plus-Function Claim Construction Analysis

Section 112, Paragraph 6² provides that a structure may be claimed as a “means...for performing a specified function.” *Masco Corp. v. United States*, 303 F.3d 1316, 1326 (Fed. Cir. 2002). If the claim language does not include the word “means” then there is a rebuttable presumption that § 112, ¶ 6 does not apply. *Williamson v. Citrix Online, LLC*, 792 F.3d 1339, 1348 (Fed. Cir. 2015). To overcome this presumption, it must be shown that the term at issue fails to “recite sufficiently definite structure” or “function without reciting sufficient structure for performing that function.” *Watts v. XL Sys., Inc.*, 232 F.3d 877, 880 (Fed. Cir. 2000).

C. Indefiniteness

The Patent Act requires claims to particularly point out and distinctly claim the subject matter regarded as the inventions. 35 U.S.C. § 112, ¶ 2. To satisfy this requirement, the claim must be read in light of the intrinsic evidence to determine whether it informs one of skill in the

² The Pre-AIA version of § 112, ¶ 6 applies to the claims in this case.

art at the time of the invention “about the scope of the invention with reasonable certainty.”

Nautilus, Inc. v. Biosig Instruments, Inc., 572 U.S. 898, 910-11 (2014). To establish that a claim is indefinite, a patent challenger must prove indefiniteness by clear and convincing evidence.

Sonix Tech. Co. v. Publ’ns Int’l, Ltd., 844 F.3d 1370, 1377 (Fed. Cir. 2017).

III. The Disputed Claim Terms

A. The ’752 Patent

The ’752 patent was issued on May 24, 2011 and claims priority to an application filed on October 23, 1998. The inventions described and claimed in the ’752 patent enable the delivery of highly reliable and customizable cloud and virtualization services to customers whose local computing platforms can be simple enough to support only a web browser.

The litigation history of the ’752 patent has particular relevance to the disputed terms now before the Court. Specifically, on July 8, 2015, IV filed a patent infringement suit against HCC Insurance Holdings, Inc. (Case No. 6:15-cv-660) (hereinafter “HCC case”) that included the ’752 patent. On August 26, 2016, Magistrate Judge Mitchell entered a Report and Recommendation construing certain terms of the ’752 patent. *See* HCC case, Dkt. 102 (hereinafter the “R&R”). Both HCC and IV filed objections to the recommended constructions, however, the case was resolved prior to a decision on the parties’ objections.

In the interest of conserving Party and Court resources, IV has adopted many of the constructions set forth in Judge Mitchell’s R&R and agrees with the reasoning supporting her conclusions. While the R&R is not determinative, it is considered persuasive. *See Maurice Mitchell Innovations, L.P. v. Intel Corp.*, 2006 WL 1751779, at *3-4 (E.D. Tex. June 21, 2006). VMware, on the other hand, agreed to adopt only one of the prior court’s constructions, and in many cases proposes the very same constructions that Magistrate Judge Mitchell rejected.

1. “exhausted” (Claims 1, 9 and 24)

IV’s Proposed Construction	VMware’s Proposed Construction
“used up to the allotted or pre-determined amount”	“unavailable for reuse”

VMware requests the Court to construe this disputed term in a manner inconsistent with the intrinsic record and contrary to the prior recommended construction from Judge Mitchell’s R&R. *See Maurice Mitchell Innovations, L.P.*, 2006 WL 1751779, at *4 (treating the prior court’s analysis as persuasive is consistent with “the premise that a uniform treatment of claim construction is desirable”).

Judge Mitchell’s logic stands true and should be followed. The intrinsic record is replete with evidence supporting IV’s proposed construction and, in contrast, nowhere unambiguously states that the disputed term is limited to being unavailable for reuse as VMware suggests. For example, the specification discloses that resource consumption by agent 22 is monitored and further consumption is halted when the amount of said resource held by agent 22 is exhausted. Ex. D at 52:31-34. The disclosure goes on to say that each agent has permission to consume up to a pre-authorized amount of each service resource when performing a task. *Id.* at 9:40-43. Further still the patent states that a service wrapper can be used to ensure that the agent 22 does not consume more than its allotted amount of any particular service resource as specified by a respective service permission. *See, e.g., id.* at 22:34-37, 22:47-50, 24:63-67, 25:1-6. One of ordinary skill in the art would understand that “exhausted” in the context of the ’752 patent means “used up to the allotted or pre-determined amount.”

IV’s construction also reflects the overall goal and novelty of the inventions embodied in the ’752 patent. The patent describes a customizable and virtualized solution that enables service provider customers to utilize the services made available by the provider in a way that avoids a “one size fits all” solution. In other words, it allows customers to pay for compute power, time

and access on an as needed basis. In this context were “exhausted” to mean “unavailable for reuse” as VMware proposes, the solution would fail at the most basic level. One customer using an amount of a service and service resource would permanently make that service and service resource unavailable to subsequent customers, even after he or she has released the resource and is no longer using it. Narrowing the term in this manner is not supported by the intrinsic record.

2. “consumed”/“consumption” (Claims 1, 9, 22 and 24)

IV’s Proposed Construction	VMware’s Proposed Construction
“used”	“used up”

IV’s proposal of construing the disputed term as “used” reflects the patentees’ intent as extensively disclosed in the specification. IV can find only a single instance where “used” is referenced alternatively as “used up.” *See* Ex. D at 8:21-23. In contrast, the patent uses “consumed” and “used” interchangeably at least half a dozen times. *See, e.g., id.* at 12:25-30 (“. . . which may also monitor the amount of each respective service resource **consumed** to ensure that no particular agent **uses** more than an amount authorized . . .”), 16:50-55 (“. . . monitor the amount of respective service resources expended, **used**, or otherwise **consumed** by one or more agents”), 25:14-16 (“at step 818, service wrapper 26 identifies the amount of each service resource actually **consumed** or **used** to execute the instruction.”) (emphasis added).

The citation in a single sentence that references “consumed” as “used up” is insufficient to support VMware’s construction, particularly in light of the widespread use of “consumed” and “used” as interchangeable. Therefore, the Court should reject VMware’s unsupported position and construe the term “consumed” as “used.”

3. “service” (Claims 1, 3, 9, 14 and 24)

IV’s Proposed Construction	VMware’s Proposed Construction
“network functionality available to agents”	“an application that is used by an agent on behalf of a principal”

The first description of the disputed term appears in the Abstract and discloses that “an agent is operable to utilize a service within the network system.” This description of the service as functionality of the network system that is available to the claimed agents is consistent throughout the entire patent. In fact, there is only a single mention of one embodiment of the patented invention where it’s noted that services may comprise a software application available to a principal which may be used by an agent on its behalf. Ex. D at 10:17-20. Presumably it is from this lone citation which VMware attempts to read into “services” the proposed limitations noted above. Such a limitation, however, is not only contrary to the basic principles of claim construction but is overwhelmingly discredited by the remaining 35 pages of description and disclosure in the ’752 patent.

For instance, the specification describes the functionality of a service as not limited to a particular form, such as an application, but rather, various network functionality made available by the operator or “service provider” to the claimed agents in order to perform an operation or task. *See, e.g., id.* at 12:14-25 (“[i]n operation, one or more agents may be set up for each user who is a subscriber to the services offered by the operator/provider of the network system . . . [e]ach agent for a particular user performs one or more tasks on behalf of that user . . . [t]o perform these tasks, each agent utilizes one or more services, during which it may consume various respective service resources.”). Examples of the claimed “services” available via the network system to the subscribers include “an e-mail service, a voice mail service, a paging/facsimile service, an address book and calendar service, and a business news and stocks information service, respectively.” *Id.* at 13:21-28, Fig. 2. Notably, these are not described as “applications” but are said to be “sub-systems” of the network system coupled to fast Ethernet hubs, and data storage and processing hardware to support the aforementioned

services. *Id.* at 13:15-20. While IV does not disagree that the disclosed services could be compiled into discrete programs, it does dispute that the specification limits the claimed services in such a way.

A further confirmation of IV's position that the claimed "services" are not limited to applications, but in fact represent functionality of the networked system, can be found in the Microsoft Computer Dictionary, p. 475 (5th ed. 2002) [hereinafter Ex. O] (emphasis added)³, which defines "services" in the context of networking as "specialized, software-based **functionality** provided by network servers." (emphasis added). This definition is consistent with the disclosures in the specification cited above and further evidences that IV's construction should be adopted.

4. Means Plus Function Elements (Claims 1, 3, 4, 6) (individually set forth in Ex. B)

As an initial matter, VMware attempts to impermissibly narrow the identified structure for each and every functional clause in Ex. B by qualifying the corresponding structure with the language "as described in" and citing cherry-picked portions of the specification thereafter. This is improper. Identifying the specific structure corresponding to the claimed function is all that is required. Any further attempt to limit the corresponding structure, for example, by limiting that structure to only a single embodiment, is not allowed. *See Micro Chem., Inc. v. Great Plains Chem. Co.*, 194 F.3d 1250, 1258 (Fed. Cir. 1999) ("Identification of corresponding structure may embrace more than the preferred embodiment.").

As one example, for the element "*means for mediating an interaction between the means for using the service and the service*," both parties identify "service wrapper 26" as

³ Although extrinsic evidence is not dispositive, contemporaneous evidence of the understanding of one of skill in the art is permitted to further support the intrinsic record. *See Phillips*, 415 F.3d at 1317-18. *See also Markman v. Westview Instruments, Inc.*, 52 F.3d 967, 980-81 (Fed. Cir. 1995).

corresponding structure. VMware, however, further limits that structure to “as described in 16:22-38.” This reads out several other instances in the specification where service wrapper 26 is describe as performing the claimed function. *See, e.g.*, Ex. D at 3:20-27, 17:43-46, 18:49-54, 25:1-24. Accordingly, the Court should disregard VMware’s narrowing pin cites.

Aside from the impermissible qualifying of corresponding structure discussed above, the parties only disagreement relates to clause 6 in Ex. B. IV and VMware are in agreement on the claimed function of clause 6, but the structure identified by VMware is too narrow and reads out broader structure clearly linked to performance of the claimed function. For example, VMware proposes that the structure is “monitor (50),” yet monitor 50 is only a sub-part of service wrapper 26—proposed by IV—that collectively performs the claimed function of “monitoring an amount of the service resource used by the network-based agent.” The specification, for instance, discloses that “service wrapper 26 identifies the amount of each service resource actually consumed or used to execute the instruction . . . [then] asks agent server to decrement the amount allotted to agent 22 by the amount actually used.” Ex. D at 25:14-18. So while IV agrees that monitor 50 plays a role in performing the claimed function (as it is a sub-component of service wrapper 26), it is clear from the above citation that service wrapper 26 as a whole is also involved. Therefore, the Court should adopt IV’s proposed structure corresponding to the function of clause 6 because it properly encompasses the entire scope of the corresponding structure.

B. The ’686 and ’726 Patents⁴

While the ’686 patent issued on December 31, 2013, and the ’726 patent issued on September 20, 2011, both claim priority to an application filed on May 11, 2000. The inventions

⁴ Because the two patents share a specification, unless otherwise noted all citations in this Section are to the ’726 patent.

described and claimed in the '686 and '726 patents relate to systems and methods for transferring virtual machines from one physical server to another physical server to help avoid outages, and thereby enable increased service continuity for business and consumer customers.

1. “modify a resource allocation” / “modified resource allocation” / “modify[ing] [the] computer resources allocated to a virtual server” ('686 patent Claims 5, 6, 7)
2. “modify[ing] a resource allocation for the virtual server” / “modifying [the] computer resources allocated to a virtual server” ('726 patent Claims 1, 4, 5 and 8)

IV's Proposed Construction	VMWare's Proposed Construction
“modif[y/ied] set of functions and features of a physical host(s) used in implementing tasks for each virtual server” / “modify[ing] a set of the functions and features of a physical host(s) used in implementing tasks for each virtual machine”	“modif[y/ied] a quality of service guarantee” / “modify[ing] [the] computer resources allocated to a virtual server”

IV's proposal for construing the disputed terms should be adopted because, (a) it correctly identifies and incorporates the intrinsic definition of “resource/resource allocation,” and (b) it accurately reflects the knowledge of one of skill in the art. VMware's proposal, on the other hand, simply reads in “quality of service guarantee” while simultaneously reading out “resource allocation.” Not only is this not helpful, but it's contrary to black letter claim construction law, and therefore, should be rejected. *See Comark Commc'ns, Inc. v. Harris Corp.*, 156 F.3d 1182, 1186-87 (Fed. Cir. 1998).

The disputed terms at issue here appear in seven claims across both the '726 and '686 patents. Since “modify” is well known and used in its plain meaning, the construction of the “resource allocation” iterations are the real dispute. IV's proposal takes its support directly from the intrinsic record which clearly and unambiguously defines “resource” as “the set of functions and features the physical host machine uses in implementing tasks for each virtual server.” Ex. F at 4:28-30. By directly incorporating the specification's definition of the disputed terms into its

proposed construction, IV's proposal captures the true scope of the disputed terms while following the well-known claim construction principals of not reading extraneous limitations into the claims. *See Comark Commc 'ns, Inc.*, 156 F.3d at 1187.

Notably, VMware reads out "resource allocation" entirely and replaces it with "quality of service guarantee," a term found only in the preamble of two of the seven claims at issue. Even were VMware arguing that the preamble of those two claims is limiting (which it is not), the preamble is not necessary to breathe life into the claims, and thus cannot be limiting. *Pitney Bowes, Inc. v. Hewlett-Packard Co.*, 182 F.3d 1298, 1305 (Fed. Cir. 1999). Accordingly, the Court should reject VMware's transparent attempt to re-write the claims and adopt IV's proposed construction for these disputed terms.

3. "resource denials" ('726 patent Claims 1, 4, 5 and 8)

IV's Proposed Construction	VMWare's Proposed Construction
"an indication that a request by the virtual server cannot be immediately serviced"	"an indication that a request by the virtual server for additional resources is either implicitly or explicitly denied"

The Court should adopt IV's proposed construction of the disputed term because "resource denials" is broadly defined in the specification and VMware's proposed construction is merely one narrow example of a single preferred embodiment. More specifically, IV's proposed construction is a direct quotation from the portion of the specification where the invention as a whole, rather than a specific embodiment, is described. Ex. F at 2:55-56 ("a resource denial **may refer to any request by the virtual server that cannot be immediately serviced**") (emphasis added). Such disclosure clearly indicates that the use of the disputed term was not intended to be limited beyond this general definition. *See Dealertrack, Inc. v. Huber*, 674 F.3d 1315, 1321-22 (Fed. Cir. 2012).

Despite black letter claim construction law directing that a term’s construction should not be limited to specific exemplary embodiments, VMware nevertheless has chosen to propose a construction that appears in the specification as an illustrative embodiment. For instance, the portion of the specification that VMware is presumably using as support for its proposal is discussing Figure 2A and notes in prefatory language that it’s “an embodiment of the overall process . . .” and only “one embodiment.” Ex. F at 5:21-34. Further down in the specification within the discussion of yet another embodiment reflected in Figure 3, VMware’s proposed construction is again noted as “an embodiment of one process for determining whether an individual resource in a virtual server has reached its resource limit . . .” and “FIG. 3 shows four examples of resource denial signals.” *Id.* at 7:41-43, 7:51-55. Accordingly, because VMware is trying to read one preferred embodiment into the disputed term in an instance where it’s otherwise clear that the disputed term is not so limited, the Court should adopt IV’s proposal.

4. “resource unavailable messages” / “denied requests to modify a resource allocation” / “resource unavailable messages resulting from denied requests to modify a resource allocation” (’686 patent Claims 5, 6, 7)

IV’s Proposed Construction	VMWare’s Proposed Construction
“an indication that a request by the virtual server cannot be immediately serviced” / “a request by the virtual server that cannot be immediately serviced”	“messages that indicate that a request to modify a resource allocation has been denied;” <i>see</i> construction of “modify a resource allocation”

As an initial matter, IV has proposed “resource unavailable messages” and “denied requests to modify a resource allocation” for construction, while VMware proposes “resource unavailable messages resulting from denied requests to modify a resource allocation.” In an effort to avoid confusion as much as possible IV will treat its two proposed terms as distinct parts of the larger term proposed by VMware.

IV’s construction stays true to the specific sequence of the claim element as a whole. As written the claims require a resource unavailable message be the result of “denied requests to

modify a resource allocation.” *See, e.g.*, Ex. E at Claim 5(a). In the summary of the invention the patentee describes how such a determination is made, disclosing that resource denials are monitored and used to determine whether the virtual server is overloaded. *See id.* at 2:62-65. The patentee then gives a general definition of “resource denials.” *See id.* at 65-66 ([a] resource denial may refer to any request by the virtual server that cannot be immediately serviced”). Thus, the disputed terms should be read as the result of a resource denial, i.e., an indication that a request by the virtual server cannot be immediately serviced.

VMware’s proposal merely rearrange the claim language. This is not only unhelpful but also completely ignores the sequence of the claim language and disclosure reflected in the intrinsic record discussed above. Therefore, since IV’s proposed construction tracks the language and sequence in the claims as well as the teachings of the specification, it should be adopted by the Court.

5. “virtual server” (’686 patent Claims 5, 6, 7); (’726 patent Claims 1-11)

IV’s Proposed Construction	VMWare’s Proposed Construction
Plain and ordinary meaning; in the alternative, “a virtual machine that resides on a physical server and uses the physical server’s resources, but has the appearance of being a separate, dedicated machine”	“a virtual server capable of receiving a quality of service guarantee from a physical host”

The disputed term “virtual server” should be given its plain and ordinary meaning. The term’s use in the specification and throughout the claims is consistent with its customary use as would be evident to one of skill in the art at the time of the invention. For instance, the background section of the specification begins by stating that the field of the invention relates to “resource allocation for a virtual server . . . and dynamically modifying the resource allocation for a virtual server. . .”). Ex. F at 1:24-27. Taking this statement at face value, if the field of art in question is resource allocation among virtual servers, one of skill in that art would readily

know what a virtual server is. Furthermore, the remainder of the background section describes the technology at issue in well-known and commonly understood terms. *See, e.g.*, Ex. F at 1:43-46 (“an ISP would prefer to offer network services to multiple customers while keeping all of the server host computers within a central location of the ISP . . .”), 1:53-55 (“most customers will neither require nor be amenable to paying for the [use] of an entire host computer”), 2:3-5 (“when servicing the needs for multiple customers having different needs, it is desirable to provide a virtual server that is dynamic, not static, in its allocation of resources”). It is clear from these disclosures that the patentee is using the term “virtual server” as it was understood in the art at the time, and therefore, should be construed according to its plain and ordinary meaning.

The file history of U.S. Pat. No. 6,985,937 (the originally issued parent of the ’726 and ’686 patents) [hereinafter “’937 patent”] similarly makes clear that the term “virtual server” is being used in its customary manner. During the prosecution of the aforementioned patent, the patentee describes his invention’s novelty over prior art virtual servers. He states that “Yu mentions that classes can be mapped to virtual servers and that these virtual servers can then be mapped to physical servers . . . however, Yu does not suggest ‘a virtual server being allocated a portion of the resources’ of a physical host.” Ex. P at File History, Oct. 5, 2004 Office Action Response, p. 14. The patentee goes on to note that “[w]hile this implies the existence of a physical server to which multiple virtual servers are mapped, Yu does not disclose, teach, or suggest allocating to a virtual server a portion of the physical server’s resources.” *Id.* at 15. In making these statements, the patentee acknowledges that virtual servers are known in the art and that he is using the term consistent with its known meaning, and not some unconventional or heretofore unknown construct he is coining a “virtual server.”

VMware's proposal, on the other hand, should be rejected. In addition to violating the oft-cited principal that terms should normally be given their plain and ordinary meaning absent clear evidence to the contrary, VMware's proposal attempts to read in a narrowing limitation that is unsupported by the intrinsic evidence. For example, VMware's proposal uses the disputed term itself ("virtual server") and simply tacks on the limitation "capable of receiving a quality of service guarantee," a phrase found in the preamble of only two of the thirteen claims at issue. As an initial matter, there is no evidence that the preamble of these two claims is limiting, and in any case, even if it were, reading the limitation of one embodiment into the construction of the disputed term is improper. *See Comark Commc 'ns, Inc.*, 156 F.3d at 1187.

Should the Court be inclined to construe the disputed term as opposed to giving it its plain and ordinary meaning, then the term should be construed as "a virtual machine that resides on a physical server and uses the physical server's resources, but has the appearance of being a separate, dedicated machine." This is supported by the intrinsic record and contemporaneous extrinsic evidence as well. The patent describes the claimed virtual servers as being "hosted" by a physical server which "dedicate[s]" portions of its physical resources to servicing the "resident" virtual servers. Ex. F at 2:42-46. Furthermore, the specification states that a customer otherwise required to purchase a full server from an ISP could avoid doing so by instead purchasing a virtual server backed by a dedicated portion of a physical host server's resources. In other words, from the customer's point of view, the virtual server functions just as an entire physical host server would. *Id.* at 1:49-57, 2:3-17. Contemporaneous dictionaries buttress the intrinsic evidence in this regard. *See* Ex. O, p. 555 (defining virtual server as "a virtual machine that resides on an HTTP server but has the appearance to the user of being a separate HTTP server"); IEEE Standard Dictionary of Electrical and Electronics Terms, p. 1182 (6th ed. 1996)

(Ex. K) (defining a virtual machine as “a functional simulation of a computer and its associated devices”). Therefore, if the Court finds that the disputed term should be construed its plain and ordinary meaning then it should be construed as IV proposes.

6. “quality of service guarantee” (’726 patent Claims 1 and 4)

IV’s Proposed Construction	VMWare’s Proposed Construction
“a guaranteed resource allotment which can be dynamically increased/modified”	“information that specifies a guaranteed amount of an assigned resource, and that can be dynamically modified” ⁵

The disputed term should not be limited to “information that specifies” or “an assigned resource,” as VMware suggests, because the intrinsic record is devoid of any such limitation. Rather, as IV proposes, the term simply refers to a guaranteed resource allotment capable of being dynamically modified. The patent describes the disputed term as requiring only two things, (a) to guarantee a resource allotment; and (b) that the resource allotment be capable of being dynamically adjusted. *See* Ex. F at 2:38-46 (“[t]he present invention **dynamically adjusts** the quality of service guarantees for virtual servers **based upon the resource demands** experienced by the virtual servers.”) (emphasis added). By inserting “information that specifies” into its proposal VMware is conflating two distinct concepts from the specification. Quality of service guarantee information is not the same thing as the disputed term. The specification discloses that **quality of service guarantee information** can be, for example, stored in a **quality of service parameter table** in each physical host. 4:47-49. The quality of service guarantee information, therefore, is merely a representation of the quality of service guarantee, not the quality of service guarantee itself. Therefore, VMware’s construction is simply incorrect.

⁵ VMware substantially changed its proposed construction of this term just hours before the parties were to file their respective briefs, prejudicing IV’s preparation of its brief.

In addition, nowhere in the specification is a quality of service guarantee described as being limited to “an assigned” resource. By inserting the term “assigned” VMware is attempting to limit the disputed term in a way not consistent with the intrinsic record.

7. “determining that a second physical host can accommodate the requested modified resource allocation” (’686 patent Claims 5, 6 and 7)

IV’s Proposed Construction	VMWare’s Proposed Construction
plain and ordinary meaning; in the alternative, “determining that a second physical host can accommodate the request(s) by the virtual server that could not be serviced immediately”	indefinite; in the alternative, “determining that a second physical host can accommodate the denied requests to modify a resource allocation”

The disputed term should be given its plain and ordinary meaning in light of the patent’s specification, which when read—even cursorily—illustrates that the disputed term is not indefinite as VMware suggests. The intrinsic record is replete with examples that indicate the disputed term is intended to have its plain and ordinary meaning. For example, the specification refers to the second physical host as physical host 160B in Figure 5. Ex. E at 4:4-8, 11:48-55. As is clear from Figure 5, physical host 160A, 160B and 160C are all identical physical host servers, a well-known construct to one of skill in the art. *Id.* at 12:6-15 (“in this example, virtual server 162B is transferred from **old physical host 160A** to **new physical host 160B.**”) (emphasis added).

The remainder of the disputed term (“can accommodate the requested modified resource allocation”) is also used in its plain and ordinary meaning. The claim requires that a request for a modified resource allocation be denied by the first physical host, followed by a determination that a second physical host can accommodate the denied request. This means exactly what one of skill in the art would think it means, that the second physical host must have enough resources to service the previously denied request. This understanding is reflected in the specification,

which notes one exemplary way to evaluate the available host resources is to perform periodic monitoring of the resource loads on each server. *Id.* at 11:12-16.

VMware's proposal that this claim term is indefinite is belied by the claim language itself and the specification as a whole. As the above citations to the intrinsic record show, there is ample evidence that the disclosure relating to the disputed term informs those skilled in the art about the scope of the invention with reasonable certainty. For instance, the resource loads of all of the physical hosts can be monitored in order to have an accurate idea of the resource requests denied by the first physical host, and the available capacity of the remaining physical hosts. *Id.* at 11:12-20. Further, the patentee describes how various embodiments would then use that information to identify a second physical host with sufficient resources. *Id.* at 11:20-35. Further still, the specification describes in detail one exemplary embodiment of the noted easiest fit heuristic, complete with exemplary algorithms. *Id.* at 11:36-57. Therefore, one of skill in the art would be well informed as to what is being claimed and VMware's arguments to the contrary fail to show otherwise and certainly fail to raise to the level of clear and convincing evidence.

VMware's alternative construction should be rejected since it basically a mirror recitation of the disputed term itself. It is also improper because VMware reads out the limitation of "the requested modified resource allocation" and reads in "the denied request to modify a resource allocation." This is unnecessary, however, because the claim language itself makes clear that the disputed "requested modified resource allocation" ultimately refers back to the denied request to modify a resource allocation:

receiving an indication that a first physical host is overloaded, wherein the indication is based on a determination that a virtual server is overloaded and wherein the determination that a virtual server is overloaded is based on one or more resource unavailable messages resulting from denied requests to modify a resource allocation;

determining that a second physical host **can accommodate the requested modified resource allocation . . .**

Id. at 14:1-9 (emphasis added). In light of this, if the disputed term needs to be construed as anything other than its plain and ordinary meaning, as VMware contends, it should be construed to incorporate the correct construction⁶ of “denied requests to modify a resource allocation,” i.e., “determining that a second physical host can accommodate the request(s) by the virtual server that could not be immediately serviced.” This is exactly the alternate construction that IV has proposed and which the Court should adopt if it is inclined not to apply the plain and ordinary meaning.

8. “determination that a virtual server is overloaded” (’686 patent Claims 5, 6 and 7)

IV’s Proposed Construction	VMWare’s Proposed Construction
plain and ordinary meaning	“determination that an average number of resource denials for a virtual server is beyond a pre-configured threshold”

The disputed term should be given its plain and ordinary meaning and not be further limited to include “an average number of resource denials . . . beyond a pre-configured threshold” as VMware suggests. Review of the intrinsic record reveals that the disputed term is used in its commonly understood manner to those of skill in the art. This is well-illustrated by the disputed term’s constituent parts themselves. First, there can be no real dispute that “determining” is used as that term is commonly understood, i.e., to ascertain, often based on research. Ex. E at 2:62-65 (“monitors resource denials received by virtual servers and determines if a virtual server is overloaded based upon the resource denials.”). Further, as discussed in depth *supra* III.B.5, “virtual server” is used as commonly understood by those of skill in the art. *See, e.g.*, Ex. E at 1:59-2:15; Ex. P at File History, Oct. 5, 2004, Office Action

⁶ Refer *supra* section III.B.4 for why this is the correct construction of “denied requests to modify a resource allocation”.

Response, p. 14-15. Finally, “overloaded” in the context of virtual computing is a well-understood construct and is being used as such in the ’686 patent’s claims and specification. *See* Ex. E at claim 5 (“the determination that a virtual server is overloaded is based on one or more resource unavailability messages”) (emphasis added).

VMware’s proposal fails to acknowledge the understanding of one of ordinary skill in the art regarding these basic building blocks of the art itself. More importantly, however, VMware’s proposal significantly limits the disputed term by importing limitations from a lone exemplary embodiment. Presumably, for support, VMware is relying on the portions of the specification which note that “in one embodiment . . . [i]f the resource denials received by a particular virtual server exceed a pre-specified limit, the virtual server is considered overloaded . . .” Ex. E at 3:1-3. Significantly, VMware fails to take into account the broader general description of the disputed term immediately preceding its selected language and within the very same exemplary embodiment. This disclosure shows all that is required with respect to the “determination that a virtual server is overloaded” is that it be based on resource denials. *Id.* at 2:62-65 (“In one embodiment, a dynamic resource configuration module monitors resource denials received by the virtual servers **and determines if a virtual server is overloaded based upon the resource denials.**”) (emphasis added). It is well settled law that “[I]t is improper to read limitations from a preferred embodiment described in the specification—even if it is the only embodiment—into the claims absent a clear indication in the intrinsic record that the patentee intended the claims to be so limited.” *Liebel-Flarsheim Co. v. Medrad, Inc.*, 358 F.3d 898, 913 (Fed. Cir. 2004).

In this instance, as discussed above, the ’686 patent’s specification and claim language do not indicate that VMware’s proposed limitations were intended by the patentee. This conclusion is additionally supported by the prosecution history of the ’686 patent, which indicates that IV’s

interpretation is indeed the full scope of disputed term as intended by the patentee. *See* Ex. E at File History, May 14, 2013 Response to Office Action, p. 5 (“[c]laims 5-7 now recite ‘the determination that a virtual server is overloaded is based on one or more resource unavailability messages resulting from denied requests to modify a resource allocation,’ which Applicants believe would be understood by a person having ordinary skill in the art.”).⁷ As discussed above *supra* III.B.3-4, a denied request to modify a resource allocation is a “resource denial.”

Therefore, it is apparent from the intrinsic record as a whole that VMware’s proposal reads in limitations from an exemplary embodiment that neither the patentee nor the patent examiner intended, thus impermissibly limits the claim. Accordingly, VMware’s proposed construction must be rejected.

9. “virtual server overload signal” (’726 patent Claims 1, 4, 5 and 8)

IV’s Proposed Construction	VMWare’s Proposed Construction
“an indication that a virtual server has been or is being denied resources”	“signal indicating that an average number of resource denials for a virtual server is beyond a pre-configured threshold”

As with its proposed construction of the preceding term—and for the same reasons—VMware’s proposal here is impermissibly limiting. IV’s proposal should be adopted because it is consistent with the claim language as well as specification disclosure. As discussed in depth *supra* III.B, the claim language surrounding the disputed term makes it clear that a virtual server overload signal is the result of a resource denial. Therefore, IV’s proposal is most consistent with the claim language and specification disclosure, whereas VMware’s proposal seeks to import limitations into the term in a way never intended by the patentee and contrary to established claim construction law.

⁷ The examiner apparently agreed with the patentee’s argument since he immediately transmitted an issue notification based on the patentee’s arguments/amendments.

10. Alleged Means Plus Function Elements ('686 patent Claim 7) ('726 patent Claims 1, 3, 5 and 7)

i. '686 patent, claim 7 “component” terms (i.e., clauses 1-3 of Ex. A)

VMware’s allegations that certain claims of the '686 patent and '726 patent are subject to 35 U.S.C. § 112 ¶ 6 must fail because: (1) “means for” is not present in any of the claims, (2) the prefatory language used does not consist of merely a nonce word followed by function, and (3) the claim language itself provides sufficient structure to avoid the application of § 112 ¶ 6.

First, VMware argues that three of the five elements, of claim 7 of the '686 patent are subject to § 112 ¶ 6. *See* Ex. A at claims 1-3. This is incorrect. The specific elements and claim overall, read in light of the patent’s specification and lack of “means for” language, illustrate that the disputed elements connote sufficient structure to avoid application of § 112 ¶ 6. *See Fisher-Rosemount Sys. v. ABB Ltd.*, 2019 WL 6830806, at *15-16 (S.D. Tex. Dec. 12, 2019).

Importantly, because the claim does not recite “means for” there is a presumption against applying § 112 ¶ 6. *See Williamson*, 792 F.3d at 1347-49. In order to overcome this presumption, the defendant must demonstrate that the claim terms fail to recite sufficiently definite structure or else recites function without reciting sufficient structure for performing that function. *See Fisher-Rosemount Sys.*, 2019 WL 6830806, at *16. Neither of these are the case with disputed elements 1-3.

VMware fails to appreciate—or chooses to ignore—the preamble and claim element immediately preceding those elements which it contends are means-plus-function. This error is fatal to VMware’s argument. For example, the preamble of claim 7 of the '686 patent states “a system for modifying the computer resources allocated **to a virtual server operating in a first physical host** of multiple physical hosts, the system comprising.” The next element recites “one

or more **processors** and one or more **memories.**” These are key for interpreting the claim as a whole and help provide structure to the later elements.

Take the recitation in the preamble of “virtual server,” and “operating in a first physical host,” for instance. These two terms are structural in nature. A virtual server is an abstraction of a physical server (i.e., a physical host), and thus both terms collectively provide at least the same structure to one of skill in the art as the term “server” alone. It is well-settled that the term “server” connotes a definite structure to one of skill in the art. *See Sound View Innovations, LLC v. Facebook, Inc.*, 2017 WL 2221177, at *10-11 (D. Del. May 19, 2017) (“[s]erver has a well-known meaning to a person of ordinary skill in the art and connotes a definite structure.”). The terms “processor” and “memory” are also well-known structural terms. *See, e.g., Fisher-Rosemount Sys.*, 2019 WL 6830806, at *16 (“processor can on its own recite at least some structure to persons of ordinary skill in the art”) (internal quotations omitted); *Advanced Mktg. Sys., LLC v. CVS Pharmacy, Inc.*, 2016 WL 1741396, at *20 (E.D. Tex. May 3, 2016) (finding processor and memory together provided sufficient structure to avoid § 112 ¶ 6). These clearly structural terms help clarify that the disputed terms’ use of “component”—generally understood to mean “a constituent part,”⁸—means the remainder of the disputed terms following “component” are parts of the earlier claimed structural elements discussed above, such as, “virtual server,” “physical host,” “processors,” and “memories.” Therefore, the claim language in light of the specification evidences that one of ordinary skill in the art would understand the claimed elements as having a specific structure sufficient to prevent the application of § 112 ¶ 6.

If the Court is inclined to find that the above discussed elements do implicate § 112 ¶ 6, IV proposes the function and structure disclosed in Exhibit A. IV’s proposed structure for

⁸ *See Netfuel, Inc. v. F5 Networks, Inc.*, 2017 WL 2834538, at *5-8 (N.D. Ill. June 29, 2017).

elements 1-3 should be adopted based on the specification's disclosures at Figure 1, 2:63-3:4, 3:59-4:8, 5:7-28 and 5:42-62, which disclose that dynamic resource configuration module 100 in combination with physical hosts 160A-C and virtual servers 162A-G determine whether a physical host is overloaded based on resource denials and the process of transferring to a second physical host.

- ii. '726 patent Claims 1, 4, 5, "virtual server resource monitor" and "program code for creating a virtual server resource monitor" terms (i.e., clause 4 of Ex. A)

VMware next asserts that the two elements in clause 4 as identified in Exhibit A are subject to § 112 ¶ 6, however, VMware again is mistaken. With respect to the "virtual server resource monitor" term, the structural disclosures in the preamble help illustrate that one of ordinary skill in the art would understand the structure associated with the disputed element. For example, "network system," "physical hosts," and "virtual server" are all structural terms known in the art as discussed *supra* III.B.10(i), and give context to the environment in which the "virtual server resource monitor" is claimed.⁹ In addition, the disputed term itself claims that the "virtual server resource monitor" is "communicatively coupled to the first physical host," two terms that provide well-known structure among those of skill in the art. *See supra* III.B.5.¹⁰

Furthermore, the specification clearly illustrates that the term "virtual server resource monitor" is not a nonce term, and in fact, has structure that would be apparent to one of skill in the art in light of the specification's disclosure. For instance, the disputed element is a component of "dynamic resource configuration module 100." Ex. F at 4:64-67. "Dynamic resource configuration module 100" is coupled to a physical host machine or resident on a

⁹ The same is true for "computer resources." *See also CSB-Sys. Int'l, Inc. v. SAP America, Inc.*, 2011 WL 3240838, at *12-14 (E.D. Pa. July 28, 2011) (noting when construing server that various dictionaries refer to it as a computer which makes network resources, such as disk drives and printers available).

¹⁰ *See also Collaborative Agreements, LLC v. Adobe Sys. Inc.*, 2015 WL 2250391, at *12-14 (W.D. Tex. May 12, 2015).

physical host and may be implemented as a software driver. *Id.* at 3:49-56. The disputed element can also be portions of the software code implementing the dynamic resource configuration module, i.e., part of a software driver. *Id.* at 4:67-5:3. Therefore, the “virtual server resource monitor” could be a portion of a software driver that is on, or coupled to, a physical server. *See id.* at Fig. 3. This is more than sufficient for one of skill in the art to know the structure associated with the disputed term and avoids § 112 ¶ 6. *See Intellectual Ventures II LLC v. Bitco Gen. Ins. Corp.*, 2016 WL 125594, at *8 (E.D. Tex. Jan. 11, 2016) (noting software can connote structure to one of skill in the art and is typically done through outline of an algorithm, flowchart or specific set of rules). Therefore, the disputed “virtual server resource monitor” term does not implicate § 112 ¶ 6 as it has sufficient structure in the claim itself and in the specification.

The second disputed term in this group is similarly situated and does not trigger application of § 112 ¶ 6 for the same reasons discussed above. Furthermore, because the second disputed term in this group includes “program code for creating,” and it has been held that “program code” and similar terms have sufficient structure so as to avoid a means-plus-function classification, this element would also be understood by one of skill in the art as connoting sufficient structure. *See RLIS, Inc. v. Allscripts Healthcare Solutions, Inc.*, 2013 WL 3772472, at *14 (S.D. Tex. July 16, 2013); *Eolas Techs., Inc. v. Adobe Sys., Inc.*, 810 F. Supp. 2d 795, 810 (E.D. Tex. 2011); *Aloft Media, LLC v. Adobe Sys., Inc.*, 570 F. Supp. 2d 887, 898 (E.D. Tex. 2008).

If the Court is inclined to find that the above discussed elements do implicate § 112 ¶ 6, IV proposes the function and structure disclosed in Exhibit A. IV’s proposed structure for element 4 should be adopted based on the specification’s disclosure at Figure 1, Figure 3, 4:64-

5:4, 5:44-65 which discloses that dynamic resource configuration module 100 includes, among other things, virtual service resource modifier, which in combination with physical hosts 160A-C monitors resource denials and sends a virtual server overload signal. VMware’s attempts to limit the structure to virtual server resource monitor 110, however, fails to account for the claimed functionality as a whole.

- iii. ’726 patent Claims 1, 4, 5, “virtual server resource modifier” and “program code for creating a virtual server resource modifier” terms (i.e., clause 5 of Ex. A)

This group of disputed terms is nearly identical to the two terms discussed immediately above except that instead of claiming a “monitor” they claim a “modifier.” For the same reasons discussed above, these two terms do not implicate § 112 ¶ 6.

First, the structural disclosures in the preamble, such as “network system,” “physical hosts,” “virtual server,” and “computer resources” are all structural terms known in the art and give the necessary context for one of ordinary skill in the art to understand the structure of the terms. *See supra* III.B.10(ii). In addition, the disputed term itself claims that the “virtual server resource modifier” is “communicatively coupled to the first physical host,” two terms that provide specific structure to those of ordinary skill in the art. *Id.*

Similar to the “monitor” terms, in this case, the specification strongly indicates that one of skill in the art would know the structure of the disputed terms. The disputed element is a component of “dynamic resource configuration module 100,” that in turn is coupled to a physical host machine or resident on a physical host and may be implemented as a software driver, of which the disputed element can be a part. Thus, just like the “virtual server resource monitor” terms, the disputed terms here could be a portion of a software driver that is on, or coupled to, a

physical server, which taken collectively are sufficient for one of skill in the art to know the structure associated with the disputed term and avoid § 112 ¶ 6. *See supra* III.B.10(ii).¹¹

The second disputed term in this group “program code for creating a virtual server resource modifier” also does not trigger application of § 112 ¶ 6 for the same reasons discussed above with respect to “virtual server resource modifier.” *Id.* As with the other “program code” elements, additional evidence that the application of means-plus-function to this element is incorrect is found in the fact that courts routinely hold “program code” and similar terms have been held to provide enough structure to avoid § 112 ¶ 6. *See RLIS, Inc.*, 2013 WL 3772472, at *14; *Eolas Techs., Inc.*, 810 F. Supp. 2d at 810; *Aloft Media, LLC*, 570 F. Supp. 2d at 898. As a result, when combined with the evidence set forth above, the disputed “program code” term would clearly be understood by one of skill in the art as connoting sufficient structure to avoid means-plus-function treatment.

If the Court is inclined to find that the above discussed elements do implicate § 112 ¶ 6, IV proposes the function and structure disclosed in Exhibit A. IV’s proposed structure for element 5 should be adopted based on the specification’s disclosure at Figure 1, Figure 2A, Figure 4, 4:64-5:20, which discloses that dynamic resource configuration module 100 includes, among other things, virtual service resource modifier, and in combination with physical hosts 160A-C and virtual servers 162A-G receives virtual server overload signals and signals a resource modification is needed. VMware’s attempts to limit the structure to virtual server resource monitor 120, however, fails to account for the claimed functionality as a whole as more than just the monitor 120 is necessary.

¹¹ *See also* Ex. F at Fig. 4 (illustrating exemplary flow chart for the disputed term).

- iv. '726 patent Claims 1, 4, 5, “a load balance[ing / er] module” and “program code for creating a load balancing module” terms (i.e., clause 6 of Ex. A)

VMware’s argument that the disputed terms “load balancing module” and “program code” for creating the same are means-plus-function elements is particularly illogical given that the term is a specific and well-known construct with a structure known by those of skill in the art. This is reflected in the plain language of the claim, the specification of the '726 patent, and as apparent from the file history.

The term itself (load balancing module) is structural. A load balancer, in the context of virtual computing is exactly what it sounds like, a module that balances the resource load across various physical hosts. As is described in the patent’s specification, this is typically implemented in software on or communicatively coupled, to one or more of the physical hosts. Ex. F at 4:67-5:3; 3:49-56. Not surprisingly, this is also how the claimed load balancing module is described in the claims themselves—communicatively coupled to the plurality of physical hosts. This is further illustrated in Figure 5, which provides an exemplary flow chart and schematic illustrating the structure of the load balancing software module, and in column 11 lines 22-52 where one embodiment of a load balancing algorithm is shown.¹² Therefore, the disputed term does not implicate § 112 ¶ 6, as it has sufficient structure in the claim itself in light of the specification.¹³

If the Court is inclined to find that the above discussed elements do implicate § 112 ¶ 6, IV proposes the function and structure disclosed in Exhibit A. IV’s proposed structure for

¹² See *Intellectual Ventures II LLC*, 2016 WL 125594, at *8 (software structure typically shown via an outline of an algorithm, flowchart or specific set of rules).

¹³ The same holds true for the second element in this group—program code for creating a load balancing module—with the addition of the established claim construction principal discussed above stating that the term “program code” and similar terms do not implicate § 112 ¶ 6. See *RLIS, Inc.*, 2013 WL 3772472, at *14; *Eolas Techs., Inc.*, 810 F. Supp. 2d at 810; *Aloft Media, LLC*, 570 F. Supp. 2d at 898.

element 6 should be adopted based on the specification's disclosure at Figure 1, Figure 5, Figure 6, 4:64-5:20, 10:53-11:20, which discloses that dynamic resource configuration module 100 includes, among other things, load balancer, and in combination with physical hosts 160A-C and virtual servers 162A-G receives virtual server resource modification signal and determines whether the physical host is overloaded and sends a host transfer signal if it is. VMware's attempts to limit the structure to physical host load balancing module 130, however, fails to account for the claimed functionality as a whole, since as seen in the above citations, more than just the load balancer 130 is required.

- v. '726 patent Claims 1, 4, 5, "dynamic virtual server mover" and "program code for creating a dynamic virtual server mover" terms (i.e., clause 7 of Ex. A)

This group of disputed terms is nearly identical to the virtual server monitor and virtual server modifier terms discussed *supra* III.B.10(ii)-10(iii), however, relate to a virtual server mover. For the same reasons discussed above in Sections 10(ii)-10(iii), these two terms do not implicate § 112 ¶ 6. Specifically, the same structural language appears in the claim and the same portions of the specification address the virtual server mover. It would be clear to one of skill in the art that the virtual server mover is described in structural terms as software coupled to or running on one or more physical host servers, which is itself structural in nature and well-understood. *See supra* III.B.10(ii)-(iii). Further, the "dynamic virtual server mover" is "communicatively coupled to the plurality of physical hosts," providing further structure. *Id.*¹⁴ Therefore, for the same reasons as discussed *supra* Section 10(ii)-10(iii), application of § 112 ¶ 6 is not required.

¹⁴ *See also* Ex. F at Fig. 6 (exemplary flow chart).

If the Court is inclined to find that the above discussed elements do implicate § 112 ¶ 6, IV proposes the function and structure disclosed in Exhibit A. IV's proposed structure for element 7 should be adopted based on the specification's disclosures at Figure 1, Figure 6, 4:64-5:20, 6:20-31, 6:39-59, 12:1-10, which disclose that dynamic resource configuration module 100 includes, among other things, dynamic virtual service mover and in combination with physical hosts 160A-C, receives the physical host transfer signal. VMware's attempts to limit the structure to dynamic virtual server mover 140, however, fails to account for the claimed functionality as a whole as more than just the mover 140 is necessary.

- vi. '726 patent Claims 3, 7, "dynamic virtual server mover further configured to" term (i.e., clause 8 of Ex. A)

The same arguments made above with respect to claims 1, 4 and 5 (*supra* Sections 10(ii), 10(iii) and 10(v)) hold true for this claim as well. The same specification support is applicable and same law indicates that this is not a means-plus-function claim. IV will not re-hash these arguments here for the sake of brevity and judicial efficiency, however, if the Court is inclined to find that this term is subject to § 112 ¶ 6, IV proposes the function and structure disclosed in Exhibit A. IV's proposed structure for element 8 should be adopted based on the same evidence and reasoning cited *supra* Section 10(v).

C. The '818 Patent

Claiming priority to December 7, 2007, the '818 patent discloses a set of inventions that enable more granular quality of service ("QoS") to be provided in the context of a virtual input/output ("I/O") environment.

- 1. "virtual [network/storage network] interface layer of an application server" (Claims 1, 17, 32, 42)

IV's Proposed Construction	VMware's Proposed Construction
Plain and ordinary meaning	"a virtual storage network interface to higher layers of the virtual node in an application server" / "virtual network interface layer to higher layers of the virtual node in an application server" / "virtual interface layer to higher layers of the virtual node in an application server"

The claim limitations "virtual [network/storage network] interface layer of an application server" should be given their plain and ordinary meaning. The term's use in the specification and throughout the claims is consistent with its customary use as would be evident to one of skill in the art at the time of the invention. For example, the specification describes examples of virtual interface layers (e.g., virtual network interface 220, virtual HBA 208a) that "emulate" layers of a networking or storage protocol stack using, for example, encapsulation techniques. *See* Ex. G at 3:67-4:62, 7:40-43, Fig. 2 (describing "protocol stack components and modules," "virtual interface drivers," and associated interface layers in exemplary embodiment). The intrinsic record thus describes virtual interface layers in a manner that is fully consistent with how they are understood in the art.

VMware's proposed inclusion of the phrase "to higher layers of the virtual node" would unnecessarily qualify the "virtual interface layer of an application server" recited in the claim. The proposed language would read out embodiments disclosed in the specification by suggesting that a virtual interface layer can interface with multiple higher layers, when in fact the specification shows that a virtual layer, in at least some instances, can interface with a single higher layer. *See, e.g., id.* at 4:14-32; 4:51-62; Fig. 2 (e.g., layer 208a interfacing with higher layer 209, layer 220 interfacing with higher layer 222, etc.).

The proposed inclusion of the additional “virtual node” qualifier is similarly extraneous and contrary to the literal expression of the claim language. Notably, the term “virtual node” (as “virtual node identifier”) appears elsewhere in claims 32 and 42, but does not appear at in claims 1 or 17, indicating that the patentee knew how to use the term in claim drafting and intentionally chose not to recite it in certain claim elements. The specification includes no limitations that would otherwise impose a “virtual node” requirement on the proposed terms. *See Cont’l Circuits LLC v. Intel Corp.*, 915 F.3d 788, 796-97 (Fed. Cir. 2019) (specification must contain expressions of manifest exclusion or restriction representing a clear disavowal to limit claim scope).

2. “hierarchical token bucket resource allocation” / “token(s)” (Claims; 1, 17, 30, 32, 33, 37-42)

IV’s Proposed Construction	VMware’s Proposed Construction
Plain and ordinary meaning	The specific class-based scheduling algorithm known in the art as the “hierarchical token bucket” / “token” as used in “hierarchical token bucket resource allocation”

The term “hierarchical token bucket resource allocation” should be given its plain and ordinary meaning, as understood by persons having ordinary skill in the art and supported by the specification. The specification explains that the claimed invention involves “a two-tier hierarchical QoS management process . . . employed in a virtual I/O server” in which “network fabric resources are allocated in a hierarchical arrangement.” *See* Ex. G at 2:3-18. Said hierarchical resource allocation can be implemented “using scheduling and queuing methods such as hierarchal token bucket.” *Id.* at 9:61-65. The specification then goes on at length to describe exactly how the claimed resource allocation scheme is achieved using hierarchical token buckets in a non-limiting manner. For example, the specification explains that “HTB includes hierarchical classes where three class types exist: root, non-leaf and leaf” and further:

HTB mechanisms allocate so-called tokens for the buckets at regular intervals. Scheduling a message or packet for transmission results in deducting an amount of tokens from a corresponding bucket, and is permitted when the corresponding bucket includes a sufficient number of tokens

See id. at 9:61-10:67, Fig. 2, Figs. 12-13. Accordingly, there is no need to construe the term “hierarchical token bucket resource allocation” beyond its plain and ordinary meaning since the intrinsic record unambiguously conveys the intended meaning of the term and does so in a manner that is entirely consistent with the well-understood meanings of the constituent terms “hierarchical token bucket” and “resource allocation.”

VMware’s proposed construction, to the extent it is ascertainable, cites to a series of extrinsic papers that teach various implementations of a hierarchical token bucket algorithm as used in the art. *See* Ex. Q (disclosing Linux implementations, a WLAN implementation, 802.11 implementation, etc.). This runs contrary to the specification’s non-limiting disclosure. Indeed, VMware’s proposal, that multiple different implementations of a hierarchical token bucket algorithm were known, supports the fact the potential applications of hierarchical token bucket are numerous and should not be limited to a specific, extrinsic definition as VMware proposes. VMware’s construction further ignores the significance of “. . . resource allocation [of bandwidth],” which qualifies the “hierarchical token bucket” language and applies it to the specific context of the claimed invention. *See, e.g., Continental Circuits, LLC*, 915 F.3d at 799 (while extrinsic evidence can shed useful light on the relevant art, it is less significant than the intrinsic record in determining legally operative meaning of disputed claim language). *See also* Ex. G, 10:15-33 (describing implementation in context of allocating bandwidth).

The standalone term “token” is similarly unambiguous and extensively discussed in the specification in a manner that is fully consistent with “hierarchical token bucket resource

allocation” and its usage in the art. For instance, “tokens” are disclosed as being used to “schedule and shape traffic . . . [where] each class or node in the hierarchy has a bucket of tokens associated with it . . . [and scheduling a packet for transmission] results in the deducting an amount of tokens from a corresponding bucket.” Ex. G at 10:23-28. Thus, because the specification discloses the use of “token” in its common parlance it should also be given its plain and ordinary meaning.

3. “enforcing . . .” / “receiv[e/ing] . . .” / “classify[ing] . . .” / “compar[e/ing] . . .” / “forward[ing] . . .” / “buffer[ing] . . .” (Claims; 1, 17, 30, 32, 37-39, 42)

IV’s Proposed Construction	VMware’s Proposed Construction
Plain and ordinary meaning	“enforcing . . . across the physical [storage network] interface of the virtual I/O server” / “receiv[e/ing] in the virtual I/O server” / “classify[ing] in the virtual I/O server” / “compar[e/ing] in the virtual I/O server” / “forward[ing] in the virtual I/O server” / “buffer[ing] in the virtual I/O server”

As an initial matter, VMware proposes to construe the leading verb of almost all the disputed terms to read in the phrase “in the virtual I/O server.” This is impermissible, however, because each of the disputed terms are well-understood computer networking terms that are used in the specification in a manner entirely consistent with their ordinary and customary usage. *See, e.g.*, Ex. G at 11:38-12:46, Fig. 5 (exemplary embodiment that identifies and discusses each of the terms at issue). The terms should be given their plain and ordinary meaning.

VMware’s attempt to insert the phrase “in the virtual I/O server” to qualify nearly every element recited in the claims at issue when the term is only found in the preamble is improper. *See Catalina Mktg. Int’l, Inc. v. Coolsavings.com, Inc.*, 289 F.3d 801, 808 (Fed. Cir. 2002) (facially structural preamble limitation did not limit claim body where not recited in body, only limited separate claim in which it appeared in claim body as well). Here, the claim drafters

defined a structurally complete invention and were well aware of the term “virtual I/O server” as evidenced by their decision to expressly include said term in the preamble of some independent claims (e.g., 32, 42) and not others (e.g., claim 1), and not to include said term in the body of any of the disputed claims.

VMware’s proposed constructions also contradict the specification and would, thus, read preferred embodiments out of the claim. The specification explains that “virtual I/O servers can be used to create logical separations between the application servers and I/O subsystems to make the I/O subsystems as logical resource units to application servers” and “allow for use of existing computing infrastructures, including hardware and software, while abstracting the operation of the intermediate I/O switch fabric.” *See, e.g.*, Ex. G at 1:31-36; 4:14-23 (explaining how “virtual I/O server” is logically, but not necessarily physically distinct from the application server(s) and I/O subsystems that it connects). By reading in a “virtual I/O server” limitation into almost every claim element, VMware is improperly attempting to redefine structural relationships between their proposed virtual I/O server limitation and other structural claim elements such as “physical [network/storage network] interface(s),” “local area network packets,” and “virtual [network/storage network] interface layer.” These proposed limitations run contrary to both the plain meaning of the claim language and the specification’s characterization of a virtual I/O server as a virtual component that is logically, but not necessarily physically, abstracted from underlying server and I/O subsystems.

4. “maintain[ing] a connection over a network fabric” (Claims 1, 17, 30, 32, 42)

IV’s Proposed Construction	VMware’s Proposed Construction
Plain and ordinary meaning	"maintaining a connection between the physical interface of the application server and the physical interface of the virtual I/O server over a network fabric"

The term “maintain[ing] a connection over a network fabric” should be given its plain and ordinary meaning. The specification makes it clear that the disputed term is being used in a manner commonly understood by those of skill in the art. For example, the specification explains that “[i]n one implementation, the virtual network interface 220 is used to access network interfaces of the virtual I/O server 106 over the I/O switch fabric, using the encapsulation module 206 to provide the interfaces to establish and maintain the connection.” Ex. G at 5:17-21; Fig. 1 (illustrating fabric 104; connections to application servers 102 and virtual I/O server 106). The intrinsic record provides no basis to deviate from the well-understood meaning of this term.

VMware’s proposed construction seeks to construe the first element of every independent claim at issue (the “maintain[ing] . . .” element) to introduce additional structural limitations that are neither expressed in the claim language nor suggested by the specification. Like the proposed constructions discussed above in Section III.C.3, VMware is attempting to insert additional structural limitations “physical interface of the application server” and “physical interface of the virtual I/O server” that are not present in the relevant claims. Such language is unnecessary and would improperly import limitations from the specification and preamble into the body of the claims. *See Continental Circuits LLC*, 915 F.3d at 796-97 (cautioning that specification must clearly and manifestly disavow claim scope to have limiting effect). *See also Retractable Techs., Inc. v. Becton, Dickinson & Co.*, 653 F.3d 1296, 1306 (Fed. Cir. 2011); *Epistar Corp. v. Int’l Trade Comm’n*, 566 F.3d 1321, 1335 (Fed. Cir. 2009). In addition, the claim drafters were well aware of the terms “virtual I/O server” (recited in preamble), “application server” (recited elsewhere in claims), and “physical [interface]” (recited elsewhere

in claims), and consciously chose not to include these known terms in the claim elements at issue.

VMware’s proposed language also runs contrary to the specification, which teaches that a virtual I/O server need not be physically distinct from the application servers and associated virtual network interface(s). *See, e.g.*, Ex. G at 4:9-13 (“virtual network interface, in one implementation, emulates an Ethernet NIC [and] plugs in at the bottom of the network stack and provides an Internet Protocol address bridged by the Virtual I/O server”). In this regard, VMware’s proposed constructions improperly seek to read additional physical, structural limitations into the claims (e.g., additional “physical interface(s)” relative to the “application server(s)” and the proposed “virtual I/O server”) that are neither literally expressed in the claim language nor required by the specification.

5. Alleged Means Plus Function Elements (Claim 17) (individually set forth in Ex. C)¹⁵

As an initial matter, VMware seeks to construe a large swath of claim 17 as subject to § 112 ¶ 6 and indefinite. However, claim 17 is not subject to § 112 ¶ 6 for at least the following reasons. First, the lack of the term “means” in claim 17 raises a presumption that the claim limitations at issue are not means-plus-function limitations. *Watts* at 881. Second, to overcome this presumption, VMware is required to demonstrate that “the claim term fails to recite sufficiently definite structure or else recites function without reciting sufficient structure for performing that function.” *Inventio AG v. ThyssenKrupp Elevator Am. Corp.*, 649 F.3d 1350, 1356 (Fed. Cir. 2011) (internal quotations/citations omitted). But the claims themselves tie any recited function to sufficiently definite structure. For example, with respect to element 1 of Ex. C, the limitation “maintain a connection, over a network fabric, to a virtual storage network

¹⁵ For ease of reference we have reproduced the parties proposed constructions of these 10 terms in Exhibit C.

interface layer of an application server, wherein the virtual storage network interface layer is associated with a virtual storage node identifier” distinctly ties the function of “maintain[ing] a connection . . . to a virtual storage network interface layer of an application server” to its structure, including a network fabric, a virtual storage network interface layer, and a virtual storage node identifier. Elements 2-3 similarly recite the limitation “physical storage network interface,” which is well-understood by those skilled in the art to be structural. *See* Ex. G at 3:4-7; Fig. 1 (denoting exemplary structure associated with “physical storage network interface”). With respect to remaining elements 4-10 of Ex. C, each disputed element either individually recites structure or relates back to the structural limitations discussed above. Ex. G at claim 17 (“connection,” “virtual storage network interface layer,” “the storage command,” “the current amount of tokens,” “destination” (elements 4-10 of Ex. C)). *See, e.g.*, Ex. G at 4:6-9; 5:10-15; 7:28-31; 11:18-22; 12:7-12; 12:41-46; 13:6-17; 16:6-19; Figs. 1-2; Figs. 4-10.

Further, claim 17 includes additional elements external to the disputed elements that convey structure such as “input/output fabric interface,” “storage network interface,” and “a memory.” These additional elements, read in conjunction with the disputed claim language and considered in their entirety indicate the presence of structure to a PHOSITA.

It is only once the established presumption is rebutted—i.e., that it has been shown that the disputed term(s) do not connote sufficiently definite structure within the claims, *and* that there’s a lack of corresponding structure within the specification—that the disputed terms could be deemed indefinite. *See Apple Inc. v. Motorola, Inc.*, 757 F.3d 1286, 1296-97 (Fed. Cir. 2014). Here, however, VMware cannot meet this burden because, as discussed above, the disputed terms include sufficient structure within the claims themselves read in light of the specification.

Further, for VMare’s argument to succeed, VMware would be required to show that “one or more input/output virtualization modules” is a nonce term that would not be understood by a PHOSITA as denoting structure, which it is not. *See Zeroclick, LLC v. Apple Inc.*, 891 F.3d 1003, 1007-09 (Fed. Cir. 2018). The term module is qualified by “input/output,” which denotes structure as understood by a PHOSITA. *Id.* (refusing to apply § 112 ¶ 6 where alleged nonce term “code” was qualified as “user interface code”). *C.f. Williamson* at 1351 (applying § 112 ¶ 6 where language “distributed learning control” was insufficient to connote structure onto phrase “distributed learning control module”).

If the Court is inclined to find that the above discussed elements do implicate § 112 ¶ 6, IV proposes the function and structure disclosed in Ex. C. As indicated therein, structure for these elements should be adopted based on the specification’s disclosure of at least 2:9-18; 3:9-11; 3:15-30; 3:43-53; 3:60-4:13; 4:27-5:32, 5:55-6:42, 7:26-67; 8:1-12:28; 12:30-13:4; 13:6-14:29; 14:65-15:23; 15:46-17:19; Figs. 1-4; and Figs. 11-13; describing the structure associated with application server 102, I/O switch fabric 104, virtual I/O server 106, HBA 108, I/O fabric interfaces 110 and 202, fabric driver stack 204, encapsulation module 206, virtual HBA 208a, and virtual network interface 220.¹⁶

D. The ’051 Patent

Claiming priority to an application filed on March 15, 2000, the ’051 patent discloses a set of inventions that enable remote service providers to provide private, overlapping address spaces to multiple customers in a virtualized server environment.

1. “customer forwarding [table(s)/information]” (claims: 1, 3)

¹⁶ See also Ex. H at Figs. 5-10 (flowcharts illustrating various implementations of the structure at issue).

IV's Proposed Construction	VMware's Proposed Construction
"table(s) containing [a set/sets] of customer specific forwarding information" / "set(s) of customer specific forwarding information"	See construction of "storing a customer lookup table, the customer lookup table storing associations between physical interfaces and tunnel identifiers identifying tunnels for private networks and a plurality of customer forwarding tables" / "storing . . . customer forwarding information, . . . the customer forwarding information associating network addresses with physical interfaces and tunnel identifiers"

IV's proposed constructions of "customer forwarding [table(s)/information]" would add the words "specific" and "set(s)" to further clarify that each individual instance of "customer forwarding table" or "customer forwarding information" should be understood as a distinct piece of information relative to the remaining "plurality" of customer forwarding tables and "multiple" customer forwarding information recited elsewhere in the claim. IV's proposed language is narrowly tailored to track both the underlying claim language and the intended meaning of the terms as expressed in the specification.

Regarding the proposed word "specific," the specification explains that "the information in the customer forwarding tables is segregated by customer because the private address spaces of different customers may overlap" Ex. H at 12:21-26 (emphasis added). *See also id.* at 11:50-11:58; 11:64-12:3; 12:60-13:7; Fig. 7; Fig. 9 (illustrating separation between sets of customer forwarding tables/information 910, 920, 930, etc.). As a result of this segregation, there is a single, "correct" customer forwarding table/information accessible to each customer from amongst the plurality of customer forwarding tables/set of multiple customer forwarding information. *Id.* at 11:52-54; 12:64-13:11; Fig. 9 (emphasis added). Thus, IV is proposing the phrase "customer specific" to help explain that each individual instance of "customer forwarding

table” and “customer forwarding information” is a distinct piece of information associated with a particular customer.

Regarding the proposed words “set” / “sets,” the above reasoning also applies insofar as each instance of “customer forwarding table” and “customer forwarding information” is defined as a logically distinct or “segregated” data structure. *Id.* The proposed phrasing “[set/sets] of customer specific forwarding information” further clarifies that the “correct” customer forwarding table/information is a logically distinct piece of information relative to the remaining instances of customer forwarding tables/information (the “plurality” / “multiple”).

In addition, the term “information” is an uncountable noun having an irregular plural form such that it is not grammatically correct to state “informations;” rather the proper plural phrasing is e.g., “pieces of information.” The proposed inclusion of the words “set/sets” clarifies the grammatical construct consistently with its usage in the claim.

VMware’s proposed constructions refer to its respective proposed constructions for the entire “storing . . .” element recited in claims 1 and 3. *See infra* III.D.5. In both instances, VMware proposes to modify the language of these disputed terms to exclude the term “customer,” and selectively exclude the term “forwarding,” where they precede and qualify the terms “table(s)” / “information.” It is unclear how removing unambiguous and important claim terms, such as “customer” and “forwarding,” serves to clarify the meaning(s) of the disputed terms in any way. Reading out meaningful claim terms and otherwise failing to provide any rationale for doing so is inconsistent with well-established principles of claim construction. *See, e.g., Bell Commc’ns. Research, Inc. v. Vitalink Commc’ns. Corp.*, 55 F.3d 615, 619-20 (Fed. Cir. 1995) (internal citations omitted) (“the language of the claim defines the scope of the protected invention . . . resort must be had in the first instance to the words of the claim”).

2. “virtual server(s)” (claims; 1, 3, 6)

IV’s Proposed Construction	VMware’s Proposed Construction
<p>Plain and ordinary meaning</p> <p>Alternatively – “virtual machine(s) that reside(s) on a physical server and use(s) the physical server’s resources, but [has/have] the appearance of being a separate, dedicated machine(s)”</p>	<p>“a process executing on a host computer that accept [sic] communications requests”</p>

The specification supports construing “virtual server” with its plain and ordinary meaning. For example, the specification explains how:

[A] server application executing on a single physical host can be programmed to process requests made to multiple network addresses. Such functionality is known as virtual hosting. In virtual hosting . . . the virtual host server can service requests to multiple network addresses or domain names. Thus, the functionality of numerous hosts is provided by a single physical host computer, servicing requests made to a plurality of network addresses and domain names by multiple customers.

Ex. H at 2:50-64 (emphasis added). This disclosure demonstrates that the patentee was using the term in a non-limiting manner.

If it is determined that “virtual server(s)” must be construed beyond its plain and ordinary meaning, IV proposes an alternate construction consistent with the specification that accounts for the context of surrounding claim language. For example, in claim 1 the term “virtual servers” is recited in the context of “a host computer containing a plurality of virtual servers which support a private network address space wherein the private network address spaces of two or more of the virtual servers overlap.” The context for claim 3 is similar. The specification explains that a “virtual host server can service requests to multiple network addresses or domain names. Thus, the functionality of numerous hosts is provided by a single physical host computer . . .” Ex. H at 2:60-64 (emphasis added). *See also id.* at Fig. 6 (illustrating logical separation/abstraction of each virtual server 660, 680 in relation to respective physical hosts 640, 650).

In other words, a virtual server is something that appears to be an ordinary server from the outside, but internally it is functionally isolated from its underlying physical host and/or other virtual machines or applications running thereon. IV's proposed alternate construction also closely tracks well established extrinsic definitions of "virtual server." *See, e.g.*, Ex. O at 555.

VMware's proposed construction fails to distinguish a virtual server from a conventional physical server in any meaningful way. For example, VMware's construction of "a process executing on a host computer that accept[s] communication requests," is inherently limiting the disputed term to "a process" running on a traditional physical server (i.e. "host computer"). Because the patentee is claiming a virtual server, i.e. a software abstraction of physical hardware operating to appear to a user as a traditional server, the specification's disclosure is broader than VMware argues. *See, e.g.*, Ex. H at 2:60-64, Fig. 6.

3. "physical interface(s)" (claims; 1, 3)

IV's Proposed Construction	VMware's Proposed Construction
Plain and ordinary meaning	"hardware that provides a point of communication between two or more devices"

The disputed term "physical interface(s)" should be given its plain and ordinary meaning. The term's use in the specification and throughout the claims is consistent with its customary use as would be evident to one of skill in the art at the time of the invention. For instance, the specification discloses "physical interface" consistent with the understanding of one of skill in the art. *See* Ex. H at 10:37-42. The intrinsic record does not otherwise limit the plain and ordinary meaning of this term.

Besides there being no need to construe the term "physical interface(s)," VMware seeks a construction that imports the additional limitation that a physical interface be "between two or more devices." For example, the specification teaches that "[a] tunnel switch comprises one or more physical interfaces" *See* Ex. H at 10:37-38. Thus, where portions of the specification

clearly make reference to a single component having one or more physical interface(s)– e.g., a tunnel switch –VMware’s proposed construction is improper because it attempts to affirmatively read in a limitation “between two or more devices,” which is not a necessary attribute of a physical interface.

4. “storing . . .” / “determining . . .” / “using . . .” (claims; 1, 3)
 - i. “storing a customer lookup table, the customer lookup table storing associations between physical interfaces and tunnel identifiers identifying tunnels for private networks and a plurality of customer forwarding tables” / “storing customer lookup information . . . the customer lookup information specifying associations between physical interfaces and tunnel identifiers identifying tunnels for private networks and multiple customer forwarding tables”

IV’s Proposed Construction	VMware’s Proposed Construction
<p>Not a claim “term” for purposes of construction;</p> <p>Alternatively – plain and ordinary meaning;</p> <p>Alternatively – “storing a customer lookup table having associations between physical interfaces and first tunnel identifiers that identify tunnels for private networks and identify a plurality of customer forwarding tables” /</p> <p>“storing customer lookup information specifying associations between physical interfaces and first tunnel identifiers identifying tunnels for private networks”</p>	<p>“storing a table that associates a customer identifier with a pair of an incoming physical interface and an incoming tunnel identifier” /</p> <p>“storing information that associates a customer identifier with a pair of an incoming physical interface and an incoming tunnel identifier”</p>

- ii. “storing a plurality of customer forwarding tables, the customer forwarding tables associating network addresses with physical interfaces and tunnel identifiers” / “storing . . . customer forwarding information . . . the customer forwarding information associating network addresses with physical interfaces and tunnel identifiers”

IV's Proposed Construction	VMware's Proposed Construction
<p>Not a claim "term" for purposes of construction;</p> <p>Alternatively – plain and ordinary meaning;</p> <p>Alternatively – “storing tables containing sets of customer specific forwarding information that each associate network addresses with physical interfaces and second tunnel identifiers” /</p> <p>“storing sets of customer specific forwarding information each associating network addresses with physical interfaces and second tunnel identifiers”</p>	<p>“storing a plurality of tables that each associate a network address with a pair of an outgoing physical interface and an outgoing tunnel identifier” /</p> <p>“storing information that associates a network address with a pair of an outgoing physical interface and an outgoing tunnel identifier”</p>

- iii. “determining the correct customer forwarding table from the customer lookup table using the physical interface and the tunnel identifier” / “determining the correct customer forwarding information from the customer lookup information using the physical interface identifier and the tunnel identifier”

IV's Proposed Construction	VMware's Proposed Construction
<p>Not a claim "term" for purposes of construction;</p> <p>Alternatively – plain and ordinary meaning;</p> <p>Alternatively – “determining the correct table with the correct set of customer specific forwarding information from the customer lookup table using the physical interface and the first tunnel identifier” /</p> <p>“determining the correct set of customer specific forwarding information from the customer lookup information using the physical interface identifier and the first tunnel identifier”</p>	<p>“determining the correct customer forwarding table from the customer lookup table using the incoming physical interface and the incoming tunnel identifier” /</p> <p>“determining the correct customer forwarding information from the customer lookup information using the incoming physical interface identifier and the incoming tunnel identifier”</p>

- iv. “determining via the customer forwarding table a physical interface and tunnel identifier associated with a network address of the transmission” / “using the customer forwarding information to identify a physical interface and tunnel identifier associated with a network address of the transmission”

IV's Proposed Construction	VMware's Proposed Construction
<p>Not a claim “term” for purposes of construction;</p> <p>Alternatively – plain and ordinary meaning;</p> <p>Alternatively – “determining, via the correct table of customer specific forwarding information, a physical interface and second tunnel identifier associated with a network address of the transmission” /</p> <p>“using the correct set of customer specific forwarding information to identify a physical interface and second tunnel identifier associated with a network address of the transmission”</p>	<p>“determining via the customer forwarding table an outgoing physical interface and an outgoing tunnel identifier associated with a network address of the transmission” /</p> <p>“using the customer forwarding information to identify an outgoing physical interface and an outgoing tunnel identifier associated with a network address of the transmission”</p>

IV does not view these entire claim elements as claim “terms” requiring a construction for the purposes of claim construction. *See, e.g., O2 Micro Int’l Ltd. v. Beyond Innovation Tech. Co.*, 521 F.3d 1351, 1362 (Fed. Cir. 2008) (district courts not required to construe every limitation present in patent's asserted claims and must only interpret scope of claim terms for which the parties have presented a fundamental dispute). *See also Acumed LLC v. Stryker Corp.*, 483 F.3d 800, 806 (Fed. Cir. 2007) (claim construction does not require purging every shred of ambiguity). However, to the extent that any of these proposed constructions are warranted, the appropriate construction is plain and ordinary meaning.

In addition, the specific constructions of key constituent claim terms found within VMware’s proposals—such as “customer forwarding [table(s)/information]” and “physical interface(s)” —are separately at issue and more appropriately focus on the specific terms. No additional claim language is ambiguous or unclear in light of the specification, which fully describes each and every element of claims 1 and 3. *See, e.g., Ex. H at 10:61-11:41; Figs. 6-7* (providing detailed walkthrough of tunnel switching embodiments as implemented in claimed private virtual networking environment).

In addition to extending beyond the usual scope of claim construction, VMware’s proposed constructions seek to insert the words “incoming,” “outgoing,” and “pair,” to qualify the terms “tunnel identifier(s)” and “physical interface(s)”. The proposed inclusion of the terms “incoming” and “outgoing” would improperly import limitations from the specification into the claim. *See Continental Circuits; Retractable Techs; Epistar Corp. supra* III.C.4. For instance, the specification contemplates the ability of a tunnel switch to switch traffic across sets of tunnels in both directions. Ex. H at 11:51-54; 12:6-9; 12:17-20; 14:20-23 (“tables 800 and 900 operate to switch transmissions in both directions” . . . “the tunnel functions as a bi-directional data pipe”) (emphasis added). Thus, contrary to VMware’s attempt to limit the disputed terms to a single directionality requirement, the patent contemplates both unidirectional and bidirectional embodiments.

VMware’s further proposed inclusion of the word “pair” to qualify the first and second elements of claim 1 and the first element of claim 3 is similarly improper and unsupported by the intrinsic record. The claim language at issue already uses the conjunctive term “and” to group “physical interface(s)” and “tunnel identifier(s)” when intended. Inserting an additional “pair” requirement, as VMware suggests, could improperly mislead a PHOSITA to believe, for example, that the claims require tunnel identifier and physical interface information to be grouped or processed in a certain manner (e.g., contemporaneously, sequentially, etc.) not required by the specification.

The specification also clarifies that “a tunnel switch comprises one or more physical interfaces, with each interface capable of carrying many multiplexed tunnels.” *Id.* at 10:37-39 (emphasis added). VMware’s proposed constructions would contradict these teachings by reading in an artificial distinction between “incoming physical interface” and “outgoing physical

interface,” which could improperly lead a PHOSITA to believe that more than one physical interface is required, contrary to the express language in the specification.

If any constructions beyond plain and ordinary meaning are necessary, IV’s proposed alternate construction would insert the words “first” and “second” to differentiate between separate instances of tunnel identifier(s), but otherwise would leave the intended meaning of the claim language intact without importing VMware’s extraneous limitations.

IV. Conclusion

For the reasons stated herein IV’s proposed constructions should be adopted.

March 6, 2020

Respectfully submitted,

/s/ Robert Gilman

DEREK GILLILAND
State Bar No. 24007239
SOREY, GILLILAND & HULL, LLP
P.O. Box 4203
109 W. Tyler St.
Longview, Texas 75601
903.212.2822 (telephone)
derek@soreylaw.com

KARL RUPP
State Bar No. 24035243
NIX PATTERSON L.L.P.
1845 Woodall Rodgers Fwy., Suite 1050
Dallas, Texas 75201
972.831.1188 (telephone)
972.444.0716 (facsimile)
krupp@nixlaw.com

OF COUNSEL:

Paul J. Hayes
phayes@princelobel.com
Matthew Vella
mvella@princelobel.com
Robert R. Gilman
rgilman@princelobel.com
Jonathan DeBlois
jdeblois@princelobel.com
Alex Breger
abreger@princelobel.com
PRINCE LOBEL TYE LLP
One International Place, Suite 3700
Boston, MA 02110
Tel: (617) 456-8000
Fax: (617) 456-8100

COUNSEL for PLAINTIFFS

CERTIFICATE OF SERVICE

I hereby certify that on the 6th day of March, 2020, I electronically filed the foregoing with the Clerk of using the CM/ECF system which will send notification of such filing to all counsel of record.

/s/ Robert Gilman